

**Exhibit 12 to Complaint
Intellectual Ventures I LLC and Intellectual Ventures II LLC**

**Example American Count VI Systems and Services
U.S. Patent No. 7,721,282 (“the ‘282 Patent”)**

The Accused Systems and Services include without limitation American systems and services that utilize Docker; all past, current, and future systems and services that operate in the same or substantially similar manner as the specifically identified systems and services; and all past, current, and future American systems and services that have the same or substantially similar features as the specifically identified systems and services (“Example American Count VI Systems and Services” or “American Systems and Services”).¹

On information and belief, the American Systems and Services use Docker in public and/or private cloud(s). For example, American posts, or has posted, job opportunities that require familiarity with Docker containerization concepts.

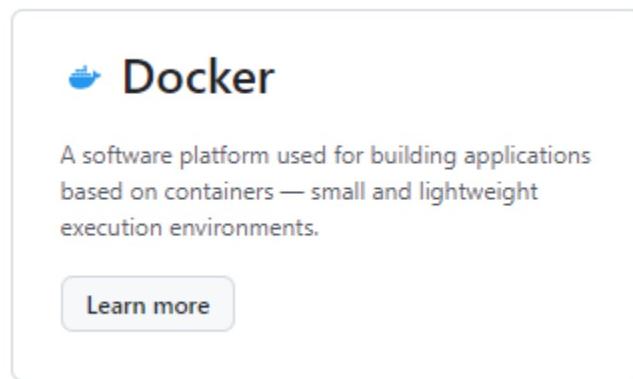
- Example of job posting for an Engineer/Sr Engineer in IT Situational Awareness at American Airlines which mentions Docker as a necessary skill for the position. https://jobs.aa.com/job/EngineerSr-Engineer%2C-IT-Situational-Awareness/75837-en_US. (Last accessed on 10/31/2024).
- Example of job posting for an Associate Developer in IT Applications at American Airlines which mentions Docker as a necessary skill for the position. https://jobs.aa.com/job/Associate-Developer%2C-IT-Applications/75816-en_US. (Last accessed on 10/31/2024).
- Example of Sr. Cloud Infra DevSecOps Engineer/Architect position at American Airlines which mentions use of Docker. <https://www.linkedin.com/in/rupa-m-b90836309/>. (Last accessed on 9/19/24).
- Example of Kubernetes Engineer position at American Airlines which mentions use of Docker. <https://www.linkedin.com/in/sridhar-pulluri-199b56250/>. (Last accessed on 9/19/24)
- Example of Software Engineer position at American Airlines which mentions use of Docker. <https://www.linkedin.com/in/pthotakura9/>. (Last accessed on 9/19/24)
- Example of Senior Developer, IT Applications position at American Airlines which mentions use of Docker. <https://www.linkedin.com/in/cj-cohorst-61a614173/>. (Last accessed on 9/19/24)

¹ Plaintiffs do not accuse the public clouds of Defendant, to the extent those services are provided by a cloud provider with a license to Plaintiffs’ patents that covers Defendant’s activities. Plaintiffs do not accuse the public clouds of Defendants if those services are provided by a cloud provider with a license to Plaintiffs’ patents that covers Defendants’ activities. Plaintiffs will produce relevant license agreements in this litigation. Plaintiffs accuse Defendant private clouds that implement Docker and non-licensed public clouds that Defendant uses to support Docker for its systems and services. Plaintiffs will provide relevant license agreements for cloud providers in discovery, to the extent any such license agreements have not already been produced. To the extent any of these licenses are relevant to Defendant’s activities, Plaintiffs will meet and confer with Defendant about the impact of such license(s). Once a protective order is entered into the case, Plaintiffs will provide further details.

- Example of Developer position at American Airlines which mentions use of Docker. <https://in.linkedin.com/in/prudhvi-kumar-rayapati>.

As another example, American has announced cloud migration of legacy technology and efforts to modernize its mainframes and servers. Source: <https://dxc.com/sg/en/insights/customer-stories/american-airlines-cloud-data-automation>. American continues to use private cloud for at least certain applications. Source: <https://www.techtarget.com/searchdatamanagement/feature/American-Airlines-lowers-data-management-costs-with-Intel>. (“American Airlines’ initial target for cost optimization was Azure Data Lake, according to Vijay Premkumar, senior manager of public and *private cloud* at the airline.”) (emphasis added).

On information and belief, other information confirmed American uses Docker technology.



Source: <https://github.com/orgs/americanair/packages>.



Top Airlines, Airports & Air Services Companies Using Docker

37,841 companies using this technology

By [Docker](#)

Docker is a software container platform. Developers use Docker to eliminate “works on my machine” problems when collaborating on code with co-workers. Operators use Docker to run and manage apps side-by-side in isolated containers to get better compute density. Enterprises use Docker to build agile software delivery pipelines. [Read less](#)



[American Airlines](#)

Technologies used by the company: 1,293

Source: <https://www.zoominfo.com/tech/23717/docker-tech-from-transportation-airline-industry-in-us-by-revenue>.

When I came to American Airlines, I tried to contribute to several InnerSource projects in our corporate VCS. One of the things I observed is that local development leaned on local installations of the application framework extensively. In the past, I've been victim of "worked on my machine" in similar setups and wanted to get a better understanding of how pervasive containers were being used. I wanted to present options on how to use containers to help remove local dependency hell. There was plenty of opportunity to leverage this approach for local development 🙌

Using containers for a local development environment can help remove impediments in situations such as:

- My stack has several dependencies that are hard to emulate:
 - Databases, and associated volumes for (short-lived) storage
 - Caches
 - Connectivity to the above, with declarative service names
- I have a bunch of application dependencies (`requirements.txt`, `packages.json`, `Gemfile`)
- There are many moving parts to reproduce a production-like environment
- My team uses different operating systems, or versions of operating systems

Windows and macOS

- Install Docker - `stable` or `edge` is fine. For the purposes of this write-up, we're using basic features.

Source: <https://tech.aa.com/2020-10-13-containers-local-dev/>.

Karl Haworth, is a Principal Staff Engineer on American Airlines Developer Experience products. His team is leveraging Chainguard Images to harden their software and strengthen the security of their software supply chain. Karl recently created an alternative, secure image for the Backstage open-source framework using Changuard's wolfi-base image. In this guest post, Karl explores the decision to use wolfi-base and the benefits he has seen in reducing vulnerabilities and shrinking the overall Backstage attack surface.

American Airlines is committed to enhancing the developer experience by implementing a frictionless self-service platform to create delightful developer experiences. By doing so, the developers can deliver value sooner to our customers. In order to achieve the stated goal, we adopted Spotify's [Backstage](#) open-source framework for accelerating the development of our internal developer portal Runway. This marks my third endeavor at establishing a developer portal which has been successful due to a heavy focus on InnerSource practices. Backstage has proven instrumental in streamlining the development process, thanks to its community-driven foundation as a solid base to build upon.

While Backstage prioritizes security, concerns arise when utilizing the base image, [node:18-bookworm-slim](#), as it contains 74 vulnerabilities ranging in severity. This has prompted us to assess the security implications when using the docker build command and [Dockerfile](#) within the Backstage framework, as the default added an additional 330 vulnerabilities. I attempted using updated node and Debian images, but they also yielded similar results, prompting me to seek out different solutions.

Source: <https://www.chainguard.dev/unchained/reducing-vulnerabilities-in-backstage-with-changuards-wolfi>. See also <https://karlhaworth.com/>.²

² Unless otherwise noted, all sources cited in this document were publicly accessible as of the filing date of the Complaint.

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
1. A system for distributing an application environment comprising:	<p>To the extent this preamble is limiting, on information and belief, the American Count VI Systems and Services are a “system for distributing an application environment.”</p> <p>Docker is an open-source platform that enables the development and distribution of applications and facilitates running applications in an isolated environment called containers. A multi-server environment capable of deploying and running Docker containers is considered a system.</p> <h2>Docker overview</h2> <p>Docker is an open platform for developing, shipping, and running applications. Docker enables you to separate your applications from your infrastructure so you can deliver software quickly. With Docker, you can manage your infrastructure in the same ways you manage your applications. By taking advantage of Docker's methodologies for shipping, testing, and deploying code, you can significantly reduce the delay between writing code and running it in production.</p> <p>Source: https://docs.docker.com/guides/docker-overview/.³</p>

³ Annotations added unless otherwise noted.

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<h2>The Docker platform</h2> <p>Docker provides the ability to package and run an application in a loosely isolated environment called a container. The isolation and security lets you run many containers simultaneously on a given host.</p> <p>Containers are lightweight and contain everything needed to run the application, so you don't need to rely on what's installed on the host. You can share containers while you work, and be sure that everyone you share with gets the same container that works in the same way.</p> <p>Docker provides tooling and a platform to manage the lifecycle of your containers:</p> <ul style="list-style-type: none">• Develop your application and its supporting components using containers.• The container becomes the unit for distributing and testing your application.• When you're ready, deploy your application into your production environment, as a container or an orchestrated service. This works the same whether your production environment is a local data center, a cloud provider, or a hybrid of the two. <p>Source: https://docs.docker.com/guides/docker-overview/.</p> <p>For example, a user can interact with Docker on the command line to run, build, pull, and/or configure containers, images, and/or volumes.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<pre> graph LR Client[Client] -- "docker run" --> DockerHost[Docker Host] Client -- "docker build" --> DockerHost Client -- "docker pull" --> DockerHost DockerHost[Docker Host] --> Images[Images] DockerHost[Docker Host] --> Containers[Containers] Images --> Containers Registry[Registry] --> Images Registry --> Containers Registry --> Extensions[Extensions] Registry --> Plugins[Plugins] </pre>
1[a] a compute node comprising a computer system;	<p>On information and belief, the American Count VI Systems and Services include “a compute node comprising a computer system.”</p> <p>Docker requires a host environment or a server to host containers in isolation from each other. Containers (e.g., a compute node) are running instances of an image, having executable packages of software, including code, runtime, system tools, libraries, and settings (e.g., comprising a computer system).</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>A container is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another. A Docker container image is a lightweight, standalone, executable package of software that includes everything needed to run an application: code, runtime, system tools, system libraries and settings.</p> <p>Source: https://www.docker.com/resources/what-container/.</p> <p style="text-align: center;">Containerized Applications</p> <pre>graph TD; subgraph CA [Containerized Applications]; direction LR; AppA[App A] --- AppB[App B] --- AppC[App C] --- AppD[App D] --- AppE[App E] --- AppF[App F]; end; CA --- Docker[Docker]; Docker --- Host[Host Operating System]; Host --- Infra[Infrastructure]</pre> <p>Source: https://www.docker.com/resources/what-container/.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>The Docker platform</p> <p>Docker provides the ability to package and run an application in a loosely isolated environment called a container. The isolation and security lets you run many containers simultaneously on a given host. Containers are lightweight and contain everything needed to run the application, so you don't need to rely on what's installed on the host. You can share containers while you work, and be sure that everyone you share with gets the same container that works in the same way.</p> <p>Source: https://docs.docker.com/get-started/docker-overview/.</p> <p>By default, a container is relatively well isolated from other containers and its host machine. You can control how isolated a container's network, storage, or other underlying subsystems are from other containers or from the host machine.</p> <p>Source: https://docs.docker.com/get-started/docker-overview/.</p> <p>For example, containers are running instances of images. Docker images include the libraries, dependencies, and other environment elements for an application to run in. The storage driver controls how the images and container instances are resident on the host.</p> <p>Images and containers</p> <p>Fundamentally, a container is nothing but a running process, with some added encapsulation features applied to it in order to keep it isolated from the host and from other containers. One of the most important aspects of container isolation is that each container interacts with its own private filesystem; this filesystem is provided by a Docker image. An image includes everything needed to run an application - the code or binary, runtimes, dependencies, and any other filesystem objects required.</p> <p>Source: https://docker-docs.uclv.cu/get-started/.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>Docker supports several storage drivers, using a pluggable architecture. The storage driver controls how images and containers are stored and managed on your Docker host. After you have read the storage driver overview, the next step is to choose the best storage driver for your workloads. Use the storage driver with the best overall performance and stability in the most usual scenarios.</p> <p>Source: https://docs.docker.com/storage/storagedriver/select-storage-driver/.</p>
1[b] a first storage unit for storing blocks of a root image of the compute node, wherein the first storage unit comprises a first non-volatile memory, wherein the root image comprises a computer program, wherein the blocks comprise sections of data, and wherein a file of the root image comprises at least one block;	<p>On information and belief, the American Count VI Systems and Services include “a first storage unit for storing blocks of a root image of the compute node, wherein the first storage unit comprises a first non-volatile memory, wherein the root image comprises a computer program, wherein the blocks comprise sections of data, and wherein a file of the root image comprises at least one block.”</p> <p>Docker containers each have their own private filesystem provided by a Docker image. Docker provides storage drivers that control how the images and containers are stored and managed on a Docker host.</p> <p>Docker supports several storage drivers, using a pluggable architecture. The storage driver controls how images and containers are stored and managed on your Docker host. After you have read the storage driver overview, the next step is to choose the best storage driver for your workloads. Use the storage driver with the best overall performance and stability in the most usual scenarios.</p> <p>Source: https://docs.docker.com/storage/storagedriver/select-storage-driver/.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>Source: https://docs.docker.com/guides/docker-overview/.</p> <p>For example, Docker stores images and containers in memory on the host system. Particularly, Docker, through its OverlayFS (Overlay File System), presents all the directories as a single consolidated directory. The lower directory is called ‘lowerdir’ and the upper directory is called ‘upperdir’. The ‘lowerdir’ is the layer that stores blocks of the base or root image. The base or root image includes the application code, dependencies, and other necessary elements for execution or runtime. Each layer, or the files therein, occupies a portion of the ‘lowerdir’.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>Image and container layers on-disk</p> <p>After downloading a five-layer image using <code>docker pull ubuntu</code>, you can see six directories under <code>/var/lib/docker/overlay2</code>.</p> <div style="background-color: #f8d7da; padding: 10px;"> <p>✖ Warning</p> <p>Don't directly manipulate any files or directories within <code>/var/lib/docker/</code>. These files and directories are managed by Docker.</p> </div> <pre>\$ ls -l /var/lib/docker/overlay2 total 24 drwx----- 5 root root 4096 Jun 20 07:36 223c2864175491657d238e2664251df13b63adb8d050924f drwx----- 3 root root 4096 Jun 20 07:36 3a36935c9df35472229c57f4a27105a136f5e4dbef0f8796 drwx----- 5 root root 4096 Jun 20 07:36 4e9fa83caff3e8f4cc83693fa407a4a9fac9573deaf48156 drwx----- 5 root root 4096 Jun 20 07:36 e8876a226237217ec61c4baf238a32992291d059fdac95ec drwx----- 5 root root 4096 Jun 20 07:36 eca1e4e1694283e001f200a667bb3cb40853cf2d1b12c29f drwx----- 2 root root 4096 Jun 20 07:36 1</pre> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p>

U.S. Patent No. 7,721,282 (Claim 1)									
Claim	Example American Count VI Systems and Services								
	<p>The diagram illustrates a Docker image structure. At the bottom is a large blue rectangle labeled "ubuntu:15.04". Above it is a smaller blue rectangle labeled "Thin R/W layer". This pattern repeats four times, with each "Thin R/W layer" pointing to the one above it. To the right of the image, there is a small icon of a padlock inside a box, indicating that the image is immutable.</p> <table border="1"><tr><td>91e54dfb1179</td><td>0 B</td></tr><tr><td>d74508fb6632</td><td>1.895 KB</td></tr><tr><td>c22013c84729</td><td>194.5 KB</td></tr><tr><td>d3a1f33e8a5a</td><td>188.1 MB</td></tr></table> <p>Source: https://docs.docker.com/storage/storagedriver/.</p> <p>As shown in the evidence below, each layer can include multiple files. For example, an image layer can include three files, where each of these files includes at least one block of data. A user can see a consolidated view via the merged directory, which is a combined view of the 'lowerdir' and 'upperdir' via overlay2, the storage driver in OverlayFS.</p>	91e54dfb1179	0 B	d74508fb6632	1.895 KB	c22013c84729	194.5 KB	d3a1f33e8a5a	188.1 MB
91e54dfb1179	0 B								
d74508fb6632	1.895 KB								
c22013c84729	194.5 KB								
d3a1f33e8a5a	188.1 MB								

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>The following diagram shows how a Docker image and a Docker container are layered. The image layer is the <code>lowerdir</code> and the container layer is the <code>upperdir</code>. If the image has multiple layers, multiple <code>lowerdir</code> directories are used. The unified view is exposed through a directory called <code>merged</code> which is effectively the containers mount point.</p> <pre> graph LR subgraph Docker_constructs [Docker constructs] direction TB Cm[Container mount] --- Cl[Container layer] Cl --- Il[Image layer] end subgraph OverlaysFS_constructs [OverlaysFS constructs] direction TB M["merged"] U["upperdir"] L["lowerdir"] end Il --- Cm Il --- Cl Cl --- Cm Il --- M Il --- U Il --- L U --- M L --- M L --- U L --- Cl L --- Cm L --- M L --- U L --- Cl L --- Cm </pre> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p> <h3>Image and container layers on-disk</h3> <p>After downloading a five-layer image using <code>docker pull ubuntu</code>, you can see six directories under <code>/var/lib/docker/overlay2</code>.</p> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p>
1[c] a second storage unit for storing a leaf image, the leaf image comprising new data blocks and changes to the blocks of the root image, wherein the second storage	<p>On information and belief, the American Count VI Systems and Services include “a first storage unit for storing blocks of a root image of the compute node, wherein the first storage unit comprises a first non-volatile memory, wherein the root image comprises a computer program, wherein the blocks comprise sections of data, and wherein a file of the root image comprises at least one block.”</p> <p>Docker includes a top writable layer, where all writes to a container (e.g., a write to add new data or modify existing data) are stored in the writable layer. The writable layer contains only the modified data</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
unit comprises a second non-volatile memory; and	<p>per the changes made by the respective container instance. The writable layer is stored in memory.</p> <h2>Container and layers</h2> <p>The major difference between a container and an image is the top writable layer. All writes to the container that add new or modify existing data are stored in this writable layer. When the container is deleted, the writable layer is also deleted. The underlying image remains unchanged.</p> <p>Source: https://docs.docker.com/storage/storagedriver/.</p> <p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the \$HOME/.cache directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an ubuntu:15.04 image.</p> <p>Source: https://docs.docker.com/storage/storagedriver/.</p> <p>For example, the changes, edits, or modifications made to a container will be stored in the 'upperdir' of that container. The 'upperdir' includes only the changed, edited, or modified file(s) by the container.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>The following diagram shows how a Docker image and a Docker container are layered. The image layer is the <code>lowerdir</code> and the container layer is the <code>upperdir</code>. If the image has multiple layers, multiple <code>lowerdir</code> directories are used. The unified view is exposed through a directory called <code>merged</code> which is effectively the containers mount point.</p> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>The diagram illustrates a Docker storage overlay2 stack. At the bottom is a blue rectangular box labeled "ubuntu:15.04". Above it is a white rectangular box containing four horizontal blue bars, each representing a layer. The first bar is labeled "91e54dfb1179" and "0 B". The second bar is labeled "d74508fb6632" and "1.895 KB". The third bar is labeled "c22013c84729" and "194.5 KB". The fourth bar is labeled "d3a1f33e8a5a" and "188.1 MB". Above these four bars are four red rectangular boxes, each labeled "Thin R/W layer". Arrows point from each of these four boxes to the corresponding blue bars above them. To the right of the stack is a small icon of a padlock inside a rectangle.</p> <p>Source: https://docs.docker.com/storage/storagedriver/.</p> <p>For example, below describes how overlay2 functions by comparing and calculating the differences or changes between the parent and its corresponding layer, if it was modified.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<pre>// DiffSize calculates the changes between the specified id // and its parent and returns the size in bytes of the changes // relative to its base filesystem directory. func (d *Driver) DiffSize(id, parent string) (int64, error) { if useNaiveDiff(d.home) !d.isParent(id, parent) { return d.naiveDiff.DiffSize(id, parent) } return directory.Size(context.TODO(), d.getDiffPath(id)) } // Diff produces an archive of the changes between the specified // layer and its parent layer which may be "". func (d *Driver) Diff(id, parent string) (io.ReadCloser, error) { if useNaiveDiff(d.home) !d.isParent(id, parent) { return d.naiveDiff.Diff(id, parent) } } // Changes produces a list of changes between the specified layer and its // parent layer. If parent is "", then all changes will be ADD changes. func (d *Driver) Changes(id, parent string) ([]archive.Change, error) { return d.naiveDiff.Changes(id, parent) }</pre> <p>Source: https://github.com/moby/moby/blob/master/daemon/graphdriver/overlay2/overlay.go.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
1[d] a union block device for interfacing between the compute node and the first and second storage units to distribute the application environment to the compute node, wherein the union block device comprises a driver, wherein the union block device creates the application environment by merging the blocks of the root image stored on the first storage unit with the blocks of the leaf image stored on the second storage unit; the union block device comprises a low-level driver for interfacing between the first and second storage units and the file system of the compute node; and the union block device, upon receiving a write request from the compute node for a sector X, creates an appropriate persistent mapping for sector X.	<p>On information and belief, the American Count VI Systems and Services include “a union block device for interfacing between the compute node and the first and second storage units to distribute the application environment to the compute node, wherein the union block device comprises a driver, wherein the union block device creates the application environment by merging the blocks of the root image stored on the first storage unit with the blocks of the leaf image stored on the second storage unit; the union block device comprises a low-level driver for interfacing between the first and second storage units and the file system of the compute node; and the union block device, upon receiving a write request from the compute node for a sector X, creates an appropriate persistent mapping for sector X.”</p> <p>Docker enables building, running, and sharing applications using containers. Docker provides a storage driver – overlay2, for managing and storing images for use. The overlay2 storage driver provides a mechanism for interfacing between the container and storage units.</p> <h2>How the <code>overlay2</code> driver works</h2> <p>OverlayFS layers two directories on a single Linux host and presents them as a single directory. These directories are called layers, and the unification process is referred to as a union mount. OverlayFS refers to the lower directory as <code>lowerdir</code> and the upper directory a <code>upperdir</code>. The unified view is exposed through its own directory called <code>merged</code>.</p> <p>The <code>overlay2</code> driver natively supports up to 128 lower OverlayFS layers. This capability provides better performance for layer-related Docker commands such as <code>docker build</code> and <code>docker commit</code>, and consumes fewer inodes on the backing filesystem.</p> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>The following diagram shows how a Docker image and a Docker container are layered. The image layer is the <code>lowerdir</code> and the container layer is the <code>upperdir</code>. If the image has multiple layers, multiple <code>lowerdir</code> directories are used. The unified view is exposed through a directory called <code>merged</code> which is effectively the containers mount point.</p> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p> <p>To create a container, the <code>overlay2</code> driver combines the directory representing the image's top layer plus a new directory for the container. The image's layers are the <code>lowerdirs</code> in the overlay and are read-only. The new directory for the container is the <code>upperdir</code> and is writable.</p> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>The diagram illustrates the Docker OverlayFS union filesystem. At the bottom is a blue box labeled "ubuntu:15.04". Above it are four blue boxes representing "Thin R/W layer"s. Arrows point from each "Thin R/W layer" to the corresponding layer in the stack above. The top-most layer contains the identifier "91e54dfb1179" and size "0 B". The second layer down contains "d74508fb6632" and "1.895 KB". The third layer down contains "c22013c84729" and "194.5 KB". The bottom-most layer contains "d3a1f33e8a5a" and "188.1 MB". To the right of the stack is a small icon of a padlock inside a box.</p>

Source: <https://docs.docker.com/storage/storagedriver/>.

Images and containers

Fundamentally, a container is nothing but a running process, with some added encapsulation features applied to it in order to keep it isolated from the host and from other containers. One of the most important aspects of container isolation is that each container interacts with its own private filesystem; this filesystem is provided by a Docker **image**. An image includes everything needed to run an application - the code or binary, runtimes, dependencies, and any other filesystem objects required.

Source: <https://docker-docs.uclv.cu/get-started/>.

Docker includes an OverlayFS union filesystem that includes a storage driver, overlay2. OverlayFS includes the image layer and the container layer, which are presented as a single directory in a process called a union mount. In a container filesystem, the overlay2 storage driver combines directories representing the image's read-only layer and writable layers.

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<h2>Use the OverlayFS storage driver</h2> <p>OverlayFS is a union filesystem.</p> <p>This page refers to the Linux kernel driver as <code>OverlayFS</code> and to the Docker storage driver as <code>overlay2</code>.</p> <h3>How the <code>overlay2</code> driver works</h3> <p>OverlayFS layers two directories on a single Linux host and presents them as a single directory. These directories are called layers, and the unification process is referred to as a union mount. OverlayFS refers to the lower directory as <code>lowerdir</code> and the upper directory a <code>upperdir</code>. The unified view is exposed through its own directory called <code>merged</code>.</p> <p>The <code>overlay2</code> driver natively supports up to 128 lower OverlayFS layers. This capability provides better performance for layer-related Docker commands such as <code>docker build</code> and <code>docker commit</code>, and consumes fewer inodes on the backing filesystem.</p> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p> <p>For example, Docker refers to the base image layers as ‘lowerdir’ and the container layer as ‘upperdir,’ and the unified view is referred to as merged.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>The following diagram shows how a Docker image and a Docker container are layered. The image layer is the <code>lowerdir</code> and the container layer is the <code>upperdir</code>. If the image has multiple layers, multiple <code>lowerdir</code> directories are used. The unified view is exposed through a directory called <code>merged</code> which is effectively the containers mount point.</p> <p>Source: https://docs.docker.com/storage/storagedriver/overlayfs-driver/.</p> <p>Furthermore, all writes are made in the top writable layer (container layer / 'upperdir'). Based on information and belief, it is assumed that writes are made in a specific portion of the writable layer, particularly in the case of modifications made to the existing data stored in an already known portion of the writable layer.</p> <h2>Container and layers</h2> <p>The major difference between a container and an image is the top writable layer. All writes to the container that add new or modify existing data are stored in this writable layer. When the container is deleted, the writable layer is also deleted. The underlying image remains unchanged.</p> <p>Source: https://docs.docker.com/storage/storagedriver/.</p>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<p>The second-lowest layer, and each higher layer, contain a file called <code>lower</code>, which denotes its parent, and a directory called <code>diff</code> which contains its contents. It also contains a <code>merged</code> directory, which contains the unified contents of its parent layer and itself, and a <code>work</code> directory which is used internally by OverlayFS.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/overlayfs-driver/.</p> <p>The first time a container writes to an existing file, that file does not exist in the container (<code>upperdir</code>). The <code>overlay2</code> driver performs a <code>copy_up</code> operation to copy the file from the image (<code>lowerdir</code>) to the container (<code>upperdir</code>). The container then writes the changes to the new copy of the file in the container layer.</p> <p>However, OverlayFS works at the file level rather than the block level. This means that all OverlayFS <code>copy_up</code> operations copy the entire file, even if the file is large and only a small part of it's being modified. This can have a noticeable impact on container write performance. However, two things are worth noting:</p> <p>Source: https://docs.docker.com/engine/storage/drivers/overlayfs-driver/.</p> <p>For example, the ‘<code>lowerdir</code>’ directory can be seen by the <code>getLowerDirs</code> function to return all the files associated with the parent or base image.</p>

U.S. Patent No. 7,721,282 (Claim 1)

Claim	Example American Count VI Systems and Services
	<pre> 468 func (d *Driver) getLowerDirs(id string) ([]string, error) { 469 var lowersArray []string 470 lowers, err := os.ReadFile(path.Join(d.dir(id), lowerFile)) 471 if err == nil { 472 for _, s := range strings.Split(string(lowers), ":") { 473 lp, err := os.Readlink(path.Join(d.home, s)) 474 if err != nil { 475 return nil, err 476 } 477 lowersArray = append(lowersArray, path.Clean(path.Join(d.home, linkDir, lp))) 478 } 479 } else if !os.IsNotExist(err) { 480 return nil, err 481 } 482 return lowersArray, nil </pre> <p>Source: https://github.com/moby/moby/blob/master/daemon/graphdriver/overlay2/overlay.go.</p> <p>Furthermore, overlay2 first looks at if there have been any modifications at the writable layer. It will then merge the lower and upper directories and return that as the mount path of the container. For example, within a container, by listing out the files.</p> <pre> 508 // Get creates and mounts the required file system for the given id and returns the mount path. 509 func (d *Driver) Get(id, mountLabel string) (_ string, retErr error) { </pre>

U.S. Patent No. 7,721,282 (Claim 1)	
Claim	Example American Count VI Systems and Services
	<pre> 527 mergedDir := path.Join(dir, mergedDirName) 528 if count := d.ctr.Increment(mergedDir); count > 1 { 529 return mergedDir, nil 530 } 531 defer func() { 532 if retErr != nil { 533 if c := d.ctr.Decrement(mergedDir); c <= 0 { 534 if mntErr := unix.Unmount(mergedDir, 0); mntErr != nil { 535 logger.Errorf("error unmounting %v: %v", mergedDir, mntErr) 536 } 537 // Cleanup the created merged directory; see the comment in Put's rmdir 538 if rmErr := unix.Rmdir(mergedDir); rmErr != nil && !os.IsNotExist(rmErr) { 539 logger.Debugf("Failed to remove %s: %v: %v", id, rmErr, err) 540 } 541 } 542 }() 543 }() 565 mountData := label.FormatMountLabel(opts, mountLabel) 566 mount := unix.Mount 567 mountTarget := mergedDir </pre> <p>Source: https://github.com/moby/moby/blob/master/daemon/graphdriver/overlay2/overlay.go.</p>